EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	11	(susan near coatney).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:49
S2	6086	network adj (cache or buffer or storage)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:50
S3	884	(network or proxy) adj cache	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:50
S4	290723	server	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:50
S5	1874	port adj (id or identification)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:51
S6	2379671	switch$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:51
S7	732145	disk or disks	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:51
S8	179	ID near router	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:51
S9	269784	internet	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:51

EAST Search History

S10	49513	router	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:52
S11	7516	(fiber or fibre) adj channel	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:52
S12	1	(disk adj (ID or identification)) same (offline)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:52
S13	11572	offline	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:53
S14	19200	(storage$2 or unit$2 or device$2 or disk$2 or server$2) adj (id or identification)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:55
S15	435	S2 and S3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:55
S16	0	S15 and S5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:55
S17	174	S15 and S6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:55
S18	8	S17 and S14	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:55
S19	0	S18 and S13	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:56

EAST Search History

S20	0	S1 and S18	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/18 13:56

Mon, 17 Apr 2006, 1:06:39 PM EST

Edit an existing query or compose a new query in the Search Query Display.

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Search Query Display

[redacted]

Recent Search Queries

#1 ((disk array)<in>metadata)

#2 (RAID level<IN>metadata)

#3 (identification OR ID<IN>metadata)

#4 (storage identification<IN>metadata)

#5 (network cache<IN>metadata)

#6 (((disk array)<in>metadata)) <AND> ((RAID level<IN>metadata))

#7 ((((disk array)<in>metadata)) <AND> ((RAID level<IN>metadata))) <AND> ((network cache<IN>metadata)) <AND> ((storage identification<IN>metadata))

#8 (((((disk array)<in>metadata)) <AND> ((RAID level<IN>metadata))) <AND> ((network cache<IN>metadata)) <AND> ((storage identification<IN>metadata))) <AND> ((identification OR ID<IN>metadata))

[redacted]

Terms used
storage identification RAID offline port switch DASD

Found **1** of **175,083**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 1 of 1

Relevance scale

1 I/O reference behavior of production database workloads and the TPC benchmarks—an analysis at the logical level

Windsor W. Hsu, Alan Jay Smith, Honesty C. Young
March 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 1
Publisher: ACM Press

Full text available: pdf(5.42 MB) Additional Information: full citation, abstract, references, citings, index terms

As improvements in processor performance continue to far outpace improvements in storage performance, I/O is increasingly the bottleneck in computer systems, especially in large database systems that manage huge amoungs of data. The key to achieving good I/O performance is to thoroughly understand its characteristics. In this article we present a comprehensive analysis of the logical I/O reference behavior of the peak productiondatabase workloads from ten of the world's largest corporatio ...

Keywords: I/O, TPC benchmarks, caching, locality, prefetching, production database workloads, reference behavior, sequentiality, workload characterization

Results 1 - 1 of 1

I/O reference behavior of production database workloads and the TPC benchmarks—an analysis at the logical level

Full text Pdf (5.42 MB)

Source **ACM Transactions on Database Systems (TODS)** archive
Volume 26 , Issue 1 (March 2001) table of contents
Pages: 96 - 143
Year of Publication: 2001
ISSN:0362-5915

Authors Windsor W. Hsu Univ. of California, Berkeley, IBM Almaden Research Center, San Jose, CA
Alan Jay Smith Univ. of California, Berkeley, IBM Almaden Research Center, San Jose, CA
Honesty C. Young IBM Almaden Research Center, San Jose, CA

Publisher ACM Press New York, NY, USA

↑ ABSTRACT

As improvements in processor performance continue to far outpace improvements in storage performance, I/O is increasingly the bottleneck in computer systems, especially in large database systems that manage huge amoungs of data. The key to achieving good I/O performance is to thoroughly understand its characteristics. In this article we present a comprehensive analysis of the logical I/O reference behavior of the peak productiondatabase workloads from ten of the world's largest corporations. In particular, we focus on how these workloads respond to different techniques for caching, prefetching, and write buffering. Our findings include several broadly applicable rules of thumb that describe how effective the various I/O optimization techniques are for the production workloads. For instance, our results indicate that the buffer pool miss ratio tends to be related to the ratio of buffer pool size to data size by an inverse square root rule. A similar fourth root rule relates the write miss ratio and the ration of buffer pool size to data size. In addition, we characterize the reference characteristics of workloads similar to the Transaction Processing Performance Council (TPC) benchmarks C (TPC-C) and D(TPC-D), which are de facto standard performance measures for online transaction processing (OLTP) systems and decision support systems (DSS), respectively. Since benchmarks such as TPC-C and TPC-D can only be used effectively if their strengths and limitations are understood, a major focus of our analysis is to identify aspects of the benchmarks that stress the system differently than the production workloads. We discover that for the most part, the reference behavior of TPC-C and TPC-D fall within the range of behavior exhibited by the production workloads. However, there are some noteworthy exceptions that affect well-known I/O optimization techniques such as caching (LRU is further from the optimal for TPC-C, while there is little sharing of pages between transactions for TPC-D), prefetching (TPC-C exhibits no significant sequentiality), and write buffering (write buffering is lees effective for the TPC benchmarks). While the two TPC benchmarks generally complement one another in reflecting the characteristics of the production

workloads, there remain aspects of the real workloads that are not represented by either of the benchmarks.

↑ REFERENCES

Note: OCR errors may be found in this Reference List extracted from the full text article. ACM has opted to expose the complete List rather than only correct and linked references.

1 Alfred V. Aho , Peter J. Denning , Jeffrey D. Ullman, Principles of Optimal Page Replacement, Journal of the ACM (JACM), v.18 n.1, p.80-93, Jan. 1971

2 ATUL, C., DONALD,H.J.,SHIBAMIYA, A., LYLE,R.W.,AND WATTS, S. J. 1988. System and method for avoiding complete index traversals in sequential and almost sequential index probes. U.S. Patent 5748952. Filed May 10, 1995. Issued May 5, 1998.

3 Maurice J. Bach, The design of the UNIX operating system, Prentice-Hall, Inc., Upper Saddle River, NJ, 1986

4 Mary G. Baker , John H. Hartman , Michael D. Kupfer , Ken W. Shirriff , John K. Ousterhout, Measurements of a distributed file system, Proceedings of the thirteenth ACM symposium on Operating systems principles, p.198-212, October 13-16, 1991, Pacific Grove, California, United States

5 BELADY, L. A. 1966. A study of replacement algorithms for a virtual-storage computer. IBM Syst. J. 5, 2, 78-101.

6 Scott D. Carson , Sanjeev Setia, Analysis of the Periodic Update Write Policy for Disk Cache, IEEE Transactions on Software Engineering, v.18 n.1, p.44-54, January 1992

7 Ignacio R. Casas , Kenneth C. Sevcik, A Buffer Management Model For Use In Predicting Overall Database System Performance, Proceedings of the Fifth International Conference on Data Engineering, p.463-469, February 06-10, 1989

8 Chung-Min Chen , Nick Roussopoulos, Adaptive Database Buffer Allocation Using Query Feedback, Proceedings of the 19th International Conference on Very Large Data Bases, p.342-353, August 24-27, 1993

9 Peter M. Chen , Edward K. Lee , Garth A. Gibson , Randy H. Katz , David A. Patterson, RAID: high-performance, reliable secondary storage, ACM Computing Surveys (CSUR), v.26 n.2, p.145-185, June 1994

10 CHOU,H.T.AND DEWITT, D. J. 1985. An evaluation of buffer management strategies for relational database systems. In Proceedings of the International Conference on Very Large Data Bases (VLDB) (Stockholm, Sweden, Aug. 1985), 127-141.

11 E. F. Codd, A relational model of data for large shared data banks, Communications of the ACM, v.13 n.6, p.377-387, June 1970

12 Edward G. Coffman, Jr. , Peter J. Denning, Operating Systems Theory, Prentice Hall Professional Technical Reference, 1990

13 D. W. Cornell , Ph.S. Yu, Integration of buffer management and query optimization in relational database environment, Proceedings of the 15th international conference on Very large data bases, p.247-255, July 1989, Amsterdam, The Netherlands

14 Asit Dan , Don Towsley, An approximate analysis of the LRU and FIFO buffer replacement schemes, Proceedings of the 1990 ACM SIGMETRICS conference on Measurement and modeling of computer systems, p.143-152, April 1990, Univ. of Colorado, Boulder, Colorado, United States

15 Asit Dan , Philip S. Yu , Jen-Yao Chung, Database Access Characterization for Buffer Hit Prediction, Proceedings of the Ninth International Conference on Data Engineering, p.134-143, April 19-23, 1993

16 Joachen Doppelhammer , Thomas Höppler , Alfons Kemper , Donald Kossmann, Database performance in the real world: TPC-D and SAP R/3, Proceedings of the 1997 ACM SIGMOD international conference on Management of data, p.123-134, May 11-15, 1997, Tucson, Arizona, United States

17 Wolfgang Effelsberg , Mary E. S. Loomis, Logical, internal, and physical reference behavior in CODASYL database systems, ACM Transactions on Database Systems (TODS), v.9 n.2, p.187-213, June 1984

18 Christos Faloutsos , Raymond T. Ng , Timos K. Sellis, Predictive Load Control for Flexible Buffer Allocation, Proceedings of the 17th International Conference on Very Large Data Bases, p.265-274, September 03-06, 1991

19 R. A. Floyd , C. Schlatter Ellis, Directory Reference Patterns in Hierarchical File Systems, IEEE Transactions on Knowledge and Data Engineering, v.1 n.2, p.238-247, June 1989

20 Goetz Graefe, Query evaluation techniques for large databases, ACM Computing Surveys (CSUR), v.25 n.2, p.73-169, June 1993

21 L. M. Haas , W. Chang , G. M. Lohman , J. McPherson , P. F. Wilms , G. Lapis , B. Lindsay , H. Pirahesh , M. J. Carey , E. Shekita, Starburst Mid-Flight: As the Dust Clears, IEEE Transactions on Knowledge and Data Engineering, v.2 n.1, p.143-160, March 1990

22 Paula Hawthorn , Michael Stonebraker, Performance analysis of a relational data base management system, Proceedings of the 1979 ACM SIGMOD international conference on Management of data, May 30-June 01, 1979, Boston, Massachusetts

23 John L. Hennessy , David A. Patterson, Computer architecture (2nd ed.): a quantitative approach, Morgan Kaufmann Publishers Inc., San Francisco, CA, 1996

24 HILL, A. V. 1913. The combinations of haemoglobin with oxygen and carbon monoxide. Biochemistry J. 7, 471-480.

25 Windsor W. Hsu , Alan J Smith , Honesty C. Young, I/O Reference Behavior of Production Database Workloads and the TPC, University of California at Berkeley, Berkeley, CA, 1999

26 HSU,W.W.,SMITH,A.J.,AND YOUNG, H. C. 1999b. Results and data for 'Analysis of the I /O characteristics of production database workloads and the TPC benchmarks'. http://www. cs.berkeley.edu /~windsorh/DBChar.

27 Windsor W. Hsu , Honesty C. Young , Alan Jay Smith, Projecting the Performance of Decision Support Workloads on Systems with Smart Storage (SmartSTOR), Proceedings of the Seventh International Conference on Parallel and Distributed Systems (ICPADS'00), p.417, July 04-07, 2000

28 Windsor W. Hsu , Alan J Smith , Honesty C. Young, Analysis of the Characteristics of Production Database Workloads and, University of California at Berkeley, Berkeley, CA, 1999

29 IBM CORP. 1997a. DB2 for OS/390 V5 Installation Guide.

30 IBM CORP. 1997b. DB2 UDB V5 Administration Guide.

31 INTEL CORP. 1999. Intel extended server memory architecture (ESMA): Overcoming the 4 GB memory barrier. http://www.intel.com/procs/servers/pentiumiii/xeon/whitepapers/ESMA. htm.

32 Theodore Johnson , Dennis Shasha, 2Q: A Low Overhead High Performance Buffer Management Replacement Algorithm, Proceedings of the 20th International Conference on Very Large Data Bases, p.439-450, September 12-15, 1994

33 KEARNS,J.P.AND DEFAZIO, S. 1983. Locality of reference in hierarchical database systems. IEEE Trans. Softw. Eng. 19, 2 (March), 128-134.

34 J. P. Kearns , S. DeFazio, Diversity in database reference behavior, ACM SIGMETRICS Performance Evaluation Review, v.17 n.1, p.11-19, May 1989

35 KING, W. F. 1971. Analysis of paging algorithms. In Proceedings of the IFIP Congress (Ljubljana, Yugoslavia, Aug. 1971), 485-490.

36 Donald E. Knuth, The art of computer programming, volume 3: (2nd ed.) sorting and searching, Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, 1998

37 Scott T. Leutenegger , Daniel Dias, A modeling study of the TPC-C benchmark, Proceedings of the 1993 ACM SIGMOD international conference on Management of data, p.22-31, May 25-28, 1993, Washington, D.C., United States

38 MCNUTT, B. 1991. A simple statistical model of cache reference locality, and its application to cache planning, measurement and control. In Proceedings of the CMG (Computer Measurement Group) Conference (Nashville, TN, Dec. 1991), 203-210.

39 MCNUTT, B. 1995. MVS DASD survey: Results and trends. In Proceedings of the CMG (Computer Measurement Group) Conference (Nashville, TN, Dec. 1995), 658-667.

40 J. Menon, A performance comparison of RAID-5 and log-structured arrays, Proceedings of the Fourth IEEE International Symposium on High Performance Distributed Computing - HPDC '95, p.167, August 02-04, 1995

41 MOGUL, J. C. 1994. A better update policy. In Proceedings of the Summer 1994 USENIX Conference (Boston, MA, June 1994), 99-111.

42 C. Mohan , Don Haderle , Bruce Lindsay , Hamid Pirahesh , Peter Schwarz, ARIES: a transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging, ACM Transactions on Database Systems (TODS), v.17 n.1, p.94-162, March 1992

43 Michael N. Nelson , Brent B. Welch , John K. Ousterhout, Caching in the Sprite network file system, ACM Transactions on Computer Systems (TOCS), v.6 n.1, p.134-154, Feb. 1988

44 Raymond Ng , Christos Faloutsos , Timos Sellis, Flexible buffer allocation based on marginal gains, Proceedings of the 1991 ACM SIGMOD international conference on Management of data, p.387-396, May 29-31, 1991, Denver, Colorado, United States

45 Victor F. Nicola , Asit Dan , Daniel M. Dias, Analysis of the generalized clock buffer replacement scheme for database transaction processing, Proceedings of the 1992 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems, p.35-46, June 01-05,

1992, Newport, Rhode Island, United States

46 Elizabeth J. O'Neil , Patrick E. O'Neil , Gerhard Weikum, The LRU-K page replacement algorithm for database disk buffering, Proceedings of the 1993 ACM SIGMOD international conference on Management of data, p.297-306, May 25-28, 1993, Washington, D.C., United States

47 John K. Ousterhout , Hervé Da Costa , David Harrison , John A. Kunze , Mike Kupfer , James G. Thompson, A trace-driven analysis of the UNIX 4.2 BSD file system, Proceedings of the tenth ACM symposium on Operating systems principles, p.15-24, December 1985, Orcas Island, Washington, United States

48 William H. Press , Saul A. Teukolsky , William T. Vetterling , Brian P. Flannery, Numerical recipes in C (2nd ed.): the art of scientific computing, Cambridge University Press, New York, NY, 1992

49 RAGAZ,N.AND RODRIGUEZ-ROSELL, J. 1976. Empirical studies of storage management in a data base system. Res. Rep. RJ 1834, IBM Research Laboratory, San Jose, CA, Oct. 1976.

50 K. K. Ramakrishnan , Prabuddha Biswas , Ramakrishna Karedla, Analysis of file I/O traces in commercial computing environments, Proceedings of the 1992 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems, p.78-90, June 01-05, 1992, Newport, Rhode Island, United States

51 John T. Robinson , Murthy V. Devarakonda, Data cache management using frequency-based replacement, Proceedings of the 1990 ACM SIGMETRICS conference on Measurement and modeling of computer systems, p.134-142, April 1990, Univ. of Colorado, Boulder, Colorado, United States

52 RODRIGUEZ-ROSELL, J. 1976. Empirical data reference behavior in data base systems. IEEE Computer 9, 11 (Nov.), 9-13.

53 Mendel Rosenblum , John K. Ousterhout, The design and implementation of a log-structured file system, Proceedings of the thirteenth ACM symposium on Operating systems principles, p.1-15, October 13-16, 1991, Pacific Grove, California, United States

54 Giovanni Maria Sacco , Mario Schkolnick, A Mechanism for Managing the Buffer Pool in a Relational Database System Using the Hot Set Model, Proceedings of the 8th International Conference on Very Large Data Bases, p.257-262, September 08-10, 1982

55 Giovanni Maria Sacco , Mario Schkolnick, Buffer management in relational database systems, ACM Transactions on Database Systems (TODS), v.11 n.4, p.473-498, Dec. 1986

56 P. Griffiths Selinger , M. M. Astrahan , D. D. Chamberlin , R. A. Lorie , T. G. Price, Access path selection in a relational database management system, Proceedings of the 1979 ACM SIGMOD international conference on Management of data, May 30-June 01, 1979, Boston, Massachusetts

57 SINGHAL,V.AND SMITH, A. J. 1997. Analysis of locking behavior in three real database systems. VLDB J. 6, 1 (Jan.), 40-52. Extended version available as Tech. Rep. CSD-94-801, Computer Science Div., Univ. of California, Berkeley, CA, Apr. 1994.

58 SMITH, A. J. 1976. Analysis of the optimal, look-ahead demand paging algorithms. SIAM J. Comput. 5, 4 (Dec.), 743-757.

59 Alan Jay Smith, Sequentiality and prefetching in database systems, ACM Transactions on Database Systems (TODS), v.3 n.3, p.223-247, Sept. 1978

60 Alan J. Smith, Disk cache—miss ratio analysis and design considerations, ACM Transactions on

Computer Systems (TOCS), v.3 n.3, p.161-203, Aug. 1985

61 SMITH, A. J. 1994. Trace driven simulation in research on computer architecture and operating systems. In Proceedings of the Conference on New Directions in Simulation for Manufacturing and Communications (Tokyo, Japan, Aug. 1994), 43-49.

62 Michael Stonebraker, Operating system support for database management, Communications of the ACM, v.24 n.7, p.412-418, July 1981

63 TENG,J.Z.AND GUMAER, R. A. 1984. Managing IBM Database 2 buffers to maximize performance. IBM Syst. J. 23, 2, 211-218.

64 James G Thompson, Efficient Analysis of Caching Systems, University of California at Berkeley, Berkeley, CA, 1987

65 TPC. 1997a. TPC Benchmark TM C Standard Specification Revision 3.3.2. Transaction Processing Performance Council.

66 TPC. 1997b. TPC Benchmark TM D Standard Specification Revision 1.3.1. Transaction Processing Performance Council.

67 TPC. 1999a. TPC Benchmark TM H Standard Specification Revision 1.1.0. Transaction Processing Performance Council.

68 TPC. 1999b. TPC Benchmark TM R Standard Specification Revision 1.0.1. Transaction Processing Performance Council.

69 Thin-Fong Tsuei , Allan N. Packer , Keng-Tai Ko, Database buffer size investigation for OLTP workloads, Proceedings of the 1997 ACM SIGMOD international conference on Management of data, p.112-122, May 11-15, 1997, Tucson, Arizona, United States

70 TUEL,JR., W. G. 1976. An analysis of buffer paging in virtual storage systems. IBM J. Res. Dev. 20, 5 (Sept.), 518-520.

71 TUEL,JR., W. G. AND RODRIGUEZ-ROSELL, J. 1975. A methodology for the evaluation of data base systems. Res. Rep. RJ 1668, IBM Research Laboratory, San Jose, CA, Oct. 1975.

72 Richard A. Uhlig , Trevor N. Mudge, Trace-driven memory simulation: a survey, ACM Computing Surveys (CSUR), v.29 n.2, p.128-170, June 1997

73 A. Inkeri Verkamo, Empirical results on locality in database referencing, Proceedings of the 1985 ACM SIGMETRICS conference on Measurement and modeling of computer systems, p.49-58, August 1985, Austin, Texas, United States

74 VISHLITZKY,N.AND OFEK, Y. 1988. Sequential cache management system utilizing the establishment of a microcache and managing the contents of such according to a threshold comparison. U.S. Patent 5706467. Filed Sep 5, 1995. Issued Jan 6, 1998.

75 WELCH, B. B. 1991. Measured performance of caching in the Sprite network file system. Comput. Syst. 4, 3 (Summer), 315-342.

76 John Wilkes , Richard Golding , Carl Staelin , Tim Sullivan, The HP AutoRAID hierarchical storage system, ACM Transactions on Computer Systems (TOCS), v.14 n.1, p.108-136, Feb. 1996

77 Philip S. Yu , Douglas W. Cornell, Buffer management based on return on consumption in a multi-query environment, The VLDB Journal — The International Journal on Very Large Data Bases, v.2 n.1, p.1-38, January 1993

78 ZHOU, S., DA COSTA, H., AND SMITH, A. J. 1985. A file system tracing package for Berkeley UNIX.In Proceedings of the 10th Usenix Conference (Portland, OR, June 1985), 407-419.

79 ZIVKOV,B.T.AND SMITH, A. J. 1997. Disk cache design and performance as evaluated in large timesharing and database systems. In Proceedings of the CMG (Computer Measurement Group) Conference (Orlando, FL, Dec. 1997), 639-658.

↑ CITINGS *4*

W. Hsu , A. J. Smith, Characteristics of I/O traffic in personal computer and server workloads, IBM Systems Journal, v.42 n.2, p.347-372, April 2003

W. Hsu , A. J. Smith, The performance impact of I/O optimizations and disk improvements, IBM Journal of Research and Development, v.48 n.2, p.255-289, March 2004

Zhenmin Li , Zhifeng Chen , Yuanyuan Zhou, Mining block correlations to improve storage performance, ACM Transactions on Storage (TOS), v.1 n.2, p.213-245, May 2005

↑ ADDITIONAL RESOURCES

Technical reports, additional information and data related to the paper are available at http://www.cs.berkeley.edu/~windsorh.

↑ INDEX TERMS

Primary Classification:
C. Computer Systems Organization

Additional Classification:
D. Software
↳ **D.4** OPERATING SYSTEMS
↳ **D.4.2** Storage Management
↳ **Subjects:** Storage hierarchies

H. Information Systems
↳ **H.3** INFORMATION STORAGE AND RETRIEVAL
↳ **H.3.4** Systems and Software
↳ **Subjects:** Performance evaluation (efficiency and effectiveness)

K. Computing Milieux
↳ **K.6** MANAGEMENT OF COMPUTING AND INFORMATION SYSTEMS
↳ **K.6.2** Installation Management
↳ **Subjects:** Benchmarks

General Terms:
Algorithms, Design, Performance

Keywords:
I/O, TPC benchmarks, caching, locality, prefetching, production database workloads, reference behavior, sequentiality, workload characterization

↑ **Collaborative Colleagues:**

Windsor W. Hsu:	Ying Chen	Qingbo Zhu	
	Yongjoon Lee		
	Shauchi Ong		
	Jih Peir		
	Jih-Kwon Peir		
	Alan J Smith		
	Alan J. Smith		
	Alan Jay Smith		
	Honesty Young		
	Honesty C. Young		
Alan Jay Smith:	Forest Baskett	Jacob R. Lorch	Nathan T. Slingerland
	Allen J. Baum	Jacob Rubin Lorch	Stephen L. Squires
	David R. Ditzel	Jih-Kwon Peir	James G. Thompson
	Jeffrey D. Gee	Chris H. Perleberg	James Gordon Thompson
	John L. Hennessy	Dionisios N. Pnevmatikatos	Hervé Touati
	Mark D. Hill	Jeffrey B. Rothman	John Tse
	Mark Donald Hill	Jeffrey Blair Rothman	Honesty C. Young
	Windsor W. Hsu	Bernie Rudin	Zeke Zalcstein
	Windsor Wee Sun Hsu	Rafael H. Saavedra	Min Zhou
	Johnny K. F. Lee	Nathan Slingerland	Barbara Tockey Zivkov
Honesty C. Young:	Ming-Syan Chen	Yo Lung Lo	Alan J. Smith
	Ying Chen	Yu-Lung Lo	Alan Jay Smith
	J. R. Goodman	Yu-long Lo	James W. Stamos
	Jian-tu Hsieh	Yu-lung Lo	Arun Swami
	Windsor W. Hsu	Robert J. T. Morris	Arun N. Swami
	Kien A. Hua	Antoine N. Mourad	Kian-Lee Tan
	Chiang Lee	Andrew R. Pleszkun	Wallapak Tavanapong
	Koujuch Liou	P. B. Schechter	Philip S. Yu
	Ming-Ling Lo	Eugene J. Shekita	
	Mingling Lo	Alan J Smith	

↑ **Peer to Peer - Readers of this Article have also read:**

- Constructing reality
  Proceedings of the 11th annual international conference on Systems documentation
  Douglas A. Powell , Norman R. Ball , Mansel W. Griffiths

- M[4]: a metamodel for data preprocessing
  Proceedings of the 4th ACM international workshop on Data warehousing and OLAP
  Anca Vaduva , Jörg-Uwe Kietz , Regina Zücker

- Successful customer relationship management in financial applications (tutorial PM-1)
  Tutorial notes of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining
  Steve Gallant , Gregory Piatesky-Shapiro , Dorian Pyle

- When expert systems are wrong
  Proceedings of the 1990 ACM SIGBDP conference on Trends and directions in expert systems
  Joseph Williams

- An example of management training in expert systems: SBA loan evaluation system
  Proceedings of the 1990 ACM SIGBDP conference on Trends and directions in expert systems
  Sudesh M. Duggal , Paul R. Popovich